UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Improving the efficiency and reliability of the route discovery process in on-demand routing protocols

Permalink

<https://escholarship.org/uc/item/9r19v992>

Authors

Garcia-Luna-Aceves, J.J.
Spohn, M.A.

Publication Date

2005-03-13

Peer reviewed

Improving the Efficiency and Reliability of the Route Discovery Process in On-demand Routing Protocols

Marco Aurélio Spohn
Computer Science Department
University of California at Santa Cruz
Santa Cruz, CA 95064
Email: maspohn@cse.ucsc.edu

J.J. Garcia-Luna-Aceves
Computer Engineering Department
University of California at Santa Cruz
Santa Cruz, CA 95064
Email: jj@cse.ucsc.edu

Abstract—*Three-hop Horizon Pruning (THP)* is an algorithm for computing a *two-hop connected dominating set (TCDS)* of the network, and has been shown to be more efficient than all prior distributed broadcasting mechanisms when a TCDS is preferred over a *connected dominating set (CDS)*. However, like all other algorithms that depend on local topology information, THP is not reliable when the topology changes frequently. We describe and analyze the *Three-hop Horizon Enhanced Pruning (THEP)*, which eliminates THP's limitations. First THEP adopts a *virtual radio range (VR)* that is shorter than the physical radio range (RR), and considers as one-hop neighbors only those nodes within VR. The gap between VR and RR works as a *buffer zone* in which nodes can move without loss of connectivity. Second, upon receiving a broadcast packet, the forwarder list in the packet header is analyzed together with the current information about the local neighborhood. Based on this, a node using THEP may decide to broadcast a packet even though it has not been selected as a forwarder by the sender. We conduct extensive simulations and show that AODV-THEP attains better performance than AODV in terms of delivery ratio, control overhead, packet collisions, and end-to-end delay.

Index Terms—Route discovery, wireless ad-hoc networks, broadcasting, domination in graphs.

I. INTRODUCTION

Broadcasting operations involve the transmission of the same packet to many receivers or the entire network. However, achieving packet broadcasting using blind flooding in a network using contention-based medium access control (MAC) protocol can induce excessive contention. This effect is called the *broadcast storm* problem [1]. To reduce the impact of broadcasting signaling or data packets, the resulting mechanism must reduce the number of nodes that must attempt to forward broadcast packets, adapt to the dynamics of a mobile ad hoc network (MANET), and run in real time with only limited knowledge of the network topology.

Three-hop horizon pruning (THP) [2] is an efficient algorithm for building *two-hop connected dominating set (TCDS)* based only on local topology information (i.e., two-hop neighborhood plus the *one-hop dominating set* for each neighbor). A TCDS is a set of nodes such that every node in the network is within *two* hops from some node in the TCDS.

This work was supported in part by CNPq (Brazil), the Baskin Chair of Computer Engineering, and by the National Science Foundation under grant CNS-0435522.

In THP, every node provides its one-hop neighbors with a list specifying one or more tuples, each with the identifier of a one-hop neighbor and a bit indicating if that neighbor dominates *any* two-hop neighbor. To forward a broadcast packet, a node tries to obtain the smallest subset of *forwarders*, which are one-hop neighbors that use some of the node's two-hop neighbors to reach *any* node that is three hops away. After such a selection of forwarders, the node broadcasts its packet with a header specifying the list of forwarders, and each forwarder in turn repeats the process.

When a broadcast protocol based on neighbor information is used, it is possible to maintain fresh routes to all nodes within two hops, because every node has the two-hop neighborhood information. For example, in on demand routing protocols (e.g., AODV[3]), it is not necessary to broadcast the *route request (RREQ)* packet to every node in the network when searching for a route: it suffices to disseminate it to a TCDS of the network.

THP has been shown to improve the performance of networks with low mobility when used for broadcasting of RREQ messages in AODV [2]. However, because THP relies on correct information about the two-hop neighborhood, its performance degrades quickly with node mobility. To address this problem, we propose the *Three-hop Horizon Enhanced Pruning (THEP)*. A *virtual radio range (VR)*, shorter than the physical radio range (RR), is used for gathering information about the local neighborhood (i.e., two-hop neighborhood). Using two different transmission powers as Wu and Dai do [4] can increase interference. Hence, we use a single radio range in THEP while still managing to have a *buffer zone* in which nodes can move without compromising network connectivity.

Upon receiving a broadcast packet, the forwarder list in the packet header is analyzed together with the current information about the local neighborhood. This is done to find inconsistencies between the most up-to-date *one-hop dominating list* and the one used by the sender to compute the sender's forwarder list. Changes in the local topology may have impacted the *one-hop dominating list*. If that is the case, a node may decide to relay a broadcast packet even though it was not selected as a forwarder by the sender.

The rest of this paper is organized as follows. Section II summarizes prior work on enhancements to broadcasting of

Fig. 1. (a) In a CDS, every node is within distance one from a dominating node (gray nodes). (b) In a TCDS, within distance two.

packets in MANETs and establishes the nomenclature used to describe our approach. Section III presents a review of THP, and Section IV presents THEP. Section V shows how the route discovery process of on-demand routing protocols can be enhanced using THEP, and show simulation results comparing AODV against AODV-THEP, in which THEP is used in AODV for the processing of route requests. Section VI concludes this work.

II. BROADCASTING IN MOBILE NETWORKS

Several broadcasting techniques have been proposed that differ among each other on the heuristics applied to reduce the redundancy on broadcast transmissions. Broadcasting protocols can be categorized into four classes [5]: *Blind flooding*, in which each node broadcasts a packet to its neighbors whenever it receives the first copy of a broadcast packet; *probability-based methods*, in which a node re-broadcasts a packet with a given probability p ; *area-based methods*, in which a node broadcasts a packet based on the information about its location and the location of its neighbors; and *neighbor-information methods*, which can be further classified as *neighbor-designated method*, in which a node that transmits a packet specifies which one-hop neighbors should forward the packet, and the *self-pruning method*, in which a node receiving a broadcast transmission decides by itself about re-broadcasting or not the packet.

Williams and Camp [5] have shown that *neighbor-information* methods are preferred over other types of broadcast approaches.

Dominating Sets (DS) play a major role in deciding the forwarder list in neighbor-designated algorithms. A DS is a set of nodes such that every node in the network is either in the set or is the neighbor of a node in the set. If the graph induced by the nodes in DS is connected, we have a *connected dominating set* (CDS) (see Figure 1 (a)). The problem of determining the *minimum connected dominating set* (MCDS) is known to be NP-complete [6]. Extensive work has been done on finding good approximations for MCDS. A protocol with a constant approximation ratio of eight has been proposed by Wan et. al. [7]. However, their approach requires that a spanning tree be constructed first in order to select the dominating nodes (forwarders), and only after the tree has been constructed a broadcast can be performed. To improve the route discovery process, an approach based on determining the CDS in real time is required. Accordingly, we focus only on techniques that satisfy this basic requirement.

Lim and Kim [8] show that the MCDS problem can be reduced to the problem of building a *minimum cost flooding tree* (MCFT). Given that an optimal solution for the MCFT problem is not feasible, they propose heuristics for flooding trees, resulting in two algorithms: *self-pruning* and *dominant pruning* (DP). They show that both algorithms perform better than *blind flooding*, and that DP outperforms the simplest form of *self-pruning*.

In DP [8] the sending node decides which adjacent nodes should relay the packet. The relaying nodes are selected using a distributed CDS algorithm, and the identifiers (IDs) of the selected nodes are piggy-backed in the packet as the forwarder list. A receiving node that is requested to forward the packet again determines the forwarder list. DP uses the *greedy set cover* (GSC) algorithm to determine the forwarder list of a packet based just on partial knowledge of the network topology. GSC recursively chooses one-hop neighbors that cover the most two-hop neighbors, repeating the process until all two-hop neighbors are covered.

A few enhancements to dominant pruning have been reported recently. Lou and Wu [9] propose two enhancements to DP: *total dominant pruning* (TDP), and *partial dominant pruning* (PDP). Spohn and Garcia-Luna-Aceves [10] propose *enhanced dominant pruning* (EDP), which is applied to AODV to show its improvements compared to DP.

All distributed algorithms that rely on knowledge of the two-hop neighborhood are prone to error in the presence of mobility. And the main reason is that nodes might have inconsistent information about the neighborhood, compromising network connectivity. Wu and Dai [11] propose a solution to address the link availability problem using two transmission ranges. Information about the neighborhood and the set of forwarders is computed using a smaller radio range. And the broadcast process is performed using a larger radio range. The objective is to give nodes a *buffer zone* in which they can move without compromising local connectivity.

Our approach differs from Wu and Dai's in that we do not use two different radio ranges. Having two radio ranges increases interference, because a larger radio range means more neighbors, and more contention. We use a virtual range for computing the neighborhood and the set of forwarders. In addition to that, we modify THP to accommodate changes in the local topology. The new approach combines efficiency and reliability, performing well even in high-mobility scenarios.

III. REVIEW OF THREE-HOP HORIZON PRUNING

In DP, the forwarder list is a set of one-hop nodes such that all two-hop nodes are covered. The approach used in THP is to make the pruning process in DP more efficient by using topology information three hops away from a given node, but incurring very limited additional signaling overhead in conveying such information.

The information about the two-hop neighborhood of a node can be disseminated by means of a *neighbor protocol* that is independent of the routing protocol, or by periodically advertising the one-hop neighbor list using HELLO messages

Fig. 2. Using THP to build a TCDS. Every node in the network is at most two-hops distant from a forwarder.

as part of the routing protocol. Without loss of generality, let us assume that nodes use HELLO messages to advertise the one-hop neighbor lists of nodes.

In addition to informing its one-hop neighbors about its one-hop neighbor list, a node also communicates its *one-hop dominating nodes* (computed using DP) to its one-hop neighbors. To reduce the space required for this additional information, the dominating nodes are encoded in a bit-map format. Because a node lists all its one-hop neighbors in its HELLO message, and because the list of forwarders is a subset of the one-hop nodes, it suffices to signal (i.e., 1 bit per node) which neighbors are dominating.

The one-hop neighbor list and the *one-hop dominating list* communicated to a node by its one-hop neighbors provide the node with a three-hop horizon of how a broadcast message can be propagated to nodes that are three hops away, even though they are unknown.

Figure 2 depicts an example of applying THP, having node a as the source. Every node in the network is within distance two from a forwarder (i.e., dominating node). Because all nodes have information about the two-hop neighborhood, a *route request* (RREQ) message can be disseminated to this TCDS without any loss to the route discovery process. In the worst case, the RREQ will reach a forwarder that is two-hops distant from the destination.

IV. THREE-HOP HORIZON ENHANCED PRUNING

Because techniques such as THP rely on an accurate view of the two-hop neighborhood, it is expected not to perform well under high mobility. To tackle this problem, we propose two enhancements to THP, and name the new algorithm *Three-Hop Horizon Enhanced Pruning* (THEP).

Neighbor information is maintained using a virtual radio range (shorter than the physical radio range): rather than using two different radio ranges (as in [11], [4]), we use neighbor location, and regard as neighbors only those nodes within virtual radio range.

To cope with changes in the local topology, information provided by the forwarder list and the freshest information about the local neighborhood are used to decide if the node should broadcast the packet even though it was not selected as a forwarder by the sender.

With this two enhancements, we expect to address the lack of reliability in the presence of high mobility. As in [11], [4], the gap between the virtual and physical ranges constitutes a

Fig. 3. Example showing enhancements to THP.

TABLE I
NOTATION

N_1^i	The set of one-hop neighbors of node n_i
N_2^i	The set of two-hop neighbors of node n_i
N_i	$N_1^i \cup N_2^i$
$D_{1\text{-hop}}^j$	<i>one-hop dominating nodes</i> of node n_j . That is, $D_{1\text{-hop}}^j \subset N_1^j$ such that $(\bigcup_{k \in D_{1\text{-hop}}^j} N_1^k) = N_2^j$
$\mathcal{F}_i$	The THP forwarder list as computed by node n_i , and included in the RREQ header
$\mathcal{C}$	List of candidates to be forwarders
$\mathcal{U}[j]$	List of <i>one-hop dominating nodes</i> of node n_j (i.e., $\mathcal{U}[j] \subset D_{1\text{-hop}}^j$) that need to be covered

buffer zone where neighbors can move without incurring loss of connectivity. And because we do not apply two different power transmissions (as in [11], [4]), it does not incur on more interference.

Because we need to know if a node is within virtual radio range, we can either use node location information (provided by GPS, for instance), or estimate the distance to the node based on the signal strength of the receiving packet [12]. In the first case, the information about the node location should be piggy-backed in the HELLO message along with the neighbor list. The second option is effective, and does not add as much complexity to the system as the first one. In any case, the exact location of the node is not needed; estimating if a node is within virtual radio range suffices. For the purpose of simulations, we assume that nodes exchange their location information using periodic HELLO messages.

Figure 3 (a,b) shows an example where a node, first within virtual range (i.e., node a), moves out of virtual range but is still within radio range. In this case, even though node a is no longer a one-hop neighbor for the purpose of forwarding computations, it is still reachable. Figure 3 (c,d) shows an example where a new node (i.e., node b) moves within virtual radio range. If node S later on (i.e., after t_1) receives a broadcast packet for which it is not listed as a forwarder (i.e., $S \ni \mathcal{F}_{\text{sender}}$), node S would still broadcast the packet in case node b becomes a *one-hop dominating node*, and there is no forwarder in $\mathcal{F}_{\text{sender}}$ covering node b .

Algorithm 1 shows the pseudo-code for THEP (see Table I

for notation). If a node is listed in the forwarder list (i.e., $\mathcal{F}_S$, available from the packet header), it means that the node must forward the packet. If that is not the case, recent modifications to the local topology may have changed the *one-hop dominating list*, and it may be different from the list used by the sender to compute the forwarder list. To check that, the node computes the *one-hop dominating list* (i.e., $D_{1\text{-hop}}^i$, advertised in periodic HELLO messages and used to compute the THEP forwarder list). If there is any node in $D_{1\text{-hop}}^i$ that is not covered by at least one forwarder in $\mathcal{F}_S$, then the node should relay the packet even though it has not been selected as a forwarder by the sender. In other words, if there is no forwarder in $\mathcal{F}_S$ covering any *one-hop dominating node*, then the broadcast might not reach the segment of the network connected to these *one-hop dominating nodes*.

Algorithm 1: THEP

```

Data      :  $n_i$ ,  $S$  (sender),  $\mathcal{F}_S$  (sender's forwarder list),  $D_{1\text{-hop}}^k$  for all  $k \in N_1^i$ 
Result    :  $\mathcal{F}_i$ , the forwarder list (if empty, do not broadcast)
begin
1   $Broadcast = True$ 
2   $\mathcal{F}_i \leftarrow \emptyset$ 
3  if  $n_i \ni \mathcal{F}_S$  then
4     $Broadcast = False$ 
    /* compute the one-hop dominating list using the most up-to-date
    neighborhood information */
5     $D_{1\text{-hop}}^i \leftarrow DP(N_i)$ 
6    for  $n_k \in D_{1\text{-hop}}^i$  do
    /* If there is any one-hop dominating node that is not covered
    by a node in  $\mathcal{F}_S$  then broadcast */
7      if  $\nexists n_l \in \mathcal{F}_S \mid n_l \in N_k^i$  then
8         $Broadcast = True$ 
9        Break
10 if  $Broadcast$  then
11    $C \leftarrow N_1^i - N_1^S$ 
    /* Select for each candidate one-hop dominating nodes other than one-
    hop neighbors and the node itself */
12   for  $n_k \in C$  do
13      $\mathcal{U}[k] \leftarrow \emptyset$ 
14     for  $n_l \in D_{1\text{-hop}}^k$  do
15       if  $n_l \notin (N_1^i + n_i)$  then
16          $\mathcal{U}[k] \leftarrow \mathcal{U}[k] + \{n_l\}$ 
    /* Exclude one-hop dominating nodes covered by another candidate in
    C */
17   for  $n_k \in C$  do
18     for  $n_m \in \mathcal{U}[k]$  do
19       if  $\exists (n_l \neq n_k) \in C \mid n_m \in N_1^l$  then
20          $\mathcal{U}[k] \leftarrow \mathcal{U}[k] - n_m$ 
21         if  $\mathcal{U}[k] == \emptyset$  then
22            $C \leftarrow C - n_k$ 
    /* For every node  $n_k \in C$ , and for every  $n_m \in \mathcal{U}[k]$ , there is no
    other  $n_l \in C$  such that  $n_m \in \mathcal{U}[l]$ ; therefore, all nodes in  $C$  are
    forwarders. */
23    $\mathcal{F}_i \leftarrow C$ 
24 return  $\mathcal{F}_i$ 
end

```

V. USING THEP FOR ROUTE DISCOVERY

We now show how THEP can be applied for the dissemination of RREQ messages in on-demand routing protocols. For the purpose of discovering a route to a destination, it suffices

that the RREQ reaches those nodes with a route to the desired destination. There are two cases to consider in terms of how THEP can be used in this context.

If routes to two-hop neighbors are maintained pro-actively, then a node that is one or two hops away from the destination can reply to the RREQ directly. On the other hand, if routes to two-hop neighbors are not available pro-actively, then a RREQ can be propagated in a number of ways once it reaches a node that is two hops away from the destination. The RREQ can be relayed using the expanding ring search with TTL set to 2. Alternatively, a node can compute forwarders within the two-hop neighborhood using a dominating set technique different than THEP (e.g., DP).

We have applied THEP as the basis for deciding which nodes should broadcast RREQ messages in AODV. The new protocol is named AODV-THEP, and it is implemented in *Qualnet* [13]. To address reliability, instead of requiring *at least one* dominating node per two-hop neighbor (when computing the *one-hop dominating list*), every two-hop neighbor is covered by *at least two* one-hop nodes (if possible). This way, it increases the chance that a two-hop node receives a RREQ.

Note that AODV-THEP would incur much less overhead if it worked over a MAC protocol that exchanged the neighbor and forwarder information that is assumed to be exchanged as part of the routing protocol itself!

Routes to one-hop neighbors are kept as in standard AODV (i.e., nodes within physical radio range). Upon receiving a HELLO message, nodes update the route to the node sending the packet. The neighbor list advertised in the HELLO message contains only the neighbors within virtual radio range, and the $D_{1\text{-hop}}$ list is also computed using the virtual neighbor list.

Nodes relaying a RREQ packet, first compute the THEP forwarder list, update the RREQ header, and only then broadcast the packet. As in standard AODV, eventually the RREQ reaches a node with a valid route to the destination, or the destination itself (considering the network is connected). Because fewer nodes relay the same RREQ packet, we expect less contention and fewer collision of packets, as well as a smaller end-to-end delay, because the RREQ message propagates faster.

A. Simulation Results

The network is composed of 50 nodes spread over an area of $1500m \times 300m$. The radio model used is a 2Mbps IEEE 802.11 device with a nominal transmission range of 250m. Transmissions are omni-directional, and receptions are directional (this increases spatial reuse [14]). Traffic sources are continuous bit rate (CBR). Only 512-bytes data packets are used. The source-destination pairs are chosen randomly among the nodes in the network. Flows last for 30s in average following an exponential distribution. Source nodes keep active flows during the entire simulation time (new destinations are randomly selected as needed). During the simulation time, an average of 580 flows are initiated, and at any given time

Fig. 4. 50 nodes, 30 flows: packet delivery ratio

there are 30 active flows. Nodes begin transmitting at 50s plus an offset uniformly chosen over a 5s period to avoid synchronization in their initial transmissions. The simulation time is set to 600 seconds, and identical mobility and traffic scenarios are used for all protocols. Initially nodes are placed uniformly over a grid. Nodes move according to the random way-point model with velocities between 1 and 20m/s. Six pause times are tested: 0s (always moving), 50s, 100s, 300s, 400s, and 600s.

Experiments are repeated for 10 trials with different random number seeds, traffic endpoints, and topologies. The topology and traffic pattern are fixed using off-line generated mobility and packet generation scripts. This means that all protocols are compared having identical node mobility and traffic demands. Each data point represents the average of the 10 trials.

Four performance metrics are evaluated: (a) *Packet delivery ratio*, the ratio of the data packets delivered to the destination to those generated by the CBR sources; (b) *average end-to-end delay* for data packets, including all possible delays caused by route discovery latency, queuing at the interface, retransmission delays at the MAC layer, and propagation and transfer times; (c) *routing load*, the number of routing packets transmitted per data packet delivered to the destination, where each hop traversed by the packet is counted as one transmission; and (d) *MAC collisions*, the number of collisions detected at the MAC layer.

To evaluate the impact of the two enhancements, we run simulations for different virtual radio ranges. The following list summarizes all variants under consideration: AODV-THEP 1.0R, with virtual range and radio range the same (this way we can see the impact of the second enhancement alone); AODV-THEP 0.85R, with virtual range set to 85% of the radio range; AODV-THEP 0.75R, with virtual range set to 75% of the radio range; AODV-THP, AODV with standard THP; and AODV with and without HELLO messages.

Figure 4 presents the packet delivery ratio results. As expected, AODV-THP does not perform very well in scenarios with frequent topology changes. One of the main reasons is that it is more difficult to get an accurate view of the local topology when it changes frequently. For static networks, AODV-THP delivers around 10% more packets compared to AODV. AODV-THEP 1.0R shows the improvement due to the

Fig. 5. 50 nodes, 30 flows: average end-to-end delay

second enhancement by itself. It shows that it helps to compare any recent changes to the local topology to check if the sender is using any stale information (i.e., the last advertised *one-hop dominating list* may not include some new *one-hop dominating node*) when computing the list of forwarders (i.e., $\mathcal{F}$). AODV-THEP 1.0R starts performing better than AODV as mobility decreases (i.e., from 300s pause time and on), and it has the best results for static networks. Even though the topology actually does not change, because of transient link failures, and increased contention, the second enhancement helps to cope with transient changes to the local topology.

When mobility is present, the two proposed enhancements to THP operating together improve the performance of AODV in all circumstances. The results show that a VR of 0.85RR is better than 0.75RR for the scenarios under consideration. This means that a buffer zone of 0.15RR is enough to reach nodes moving out of the virtual range, and that it is better to keep more nodes within the VR when computing the forwarder list.

Figure 5 presents the average end-to-end delay results. Because around 580 flows are initiated during the simulation time, we observe that the large number of redundant RREQ transmissions affect the end-to-end delay in AODV. AODV incurs twice to three times as much delay than any other variant. So, it pays off pruning redundant broadcast transmissions, because it reduces contention. The two enhancements proved to be effective in reducing delay, while sustaining a high delivery ratio. The extra control overhead introduced by periodic HELLO messages in AODV does not have much impact on the end-to-end delay, because most of the routing load comes from RREQ transmissions. Together with the previous results for the delivery ratio, we can see that the reduction of redundant broadcast transmissions translate in a faster response to the route discovery process, which results in more packets being delivered at a smaller signaling cost.

Figure 6 shows the normalized routing overhead results. All the THEP variants present a much smaller overhead than AODV, because of the reduction on the number of redundant broadcast transmissions. As for the impact of the HELLO messages, we observe slightly more control overhead in AODV when HELLO messages are present. For static networks, AODV-THP presents the most cost-effective performance; its

Fig. 6. 50 nodes, 30 flows: normalized routing overhead

Fig. 7. 50 nodes, 30 flows: number of MAC collisions

delivery ratio is the second, and it has the smallest end-to-end delay and control overhead. On the other hand, AODV-THEP shows better performance than the other protocols in high mobility scenarios.

Figure 7 presents the results for the number of collisions of packets. AODV with and without HELLOs presents similar results, showing that the introduction of HELLO messages is not responsible for increasing the number of collision of packets. On the contrary, the redundancy of RREQ transmissions is the cause for more contention and collisions. For static networks, AODV-THP presents the best overall performance with the only exception of a slightly smaller delivery ratio than AODV-THEP 1.0R. With mobility, even though the new enhancements incur slightly more packet collisions, they do improve the overall performance of the network by delivering more packets, with smaller delays, and less control overhead. Because there is a clear trade-off between efficiency and reliability, the two enhancements increase the reliability at the cost of increasing the number of redundant broadcast transmissions with respect to THP.

VI. CONCLUSIONS

Three-hop Horizon Pruning (THP) has been shown to outperform the best-performing self-pruning and neighbor-designated algorithms known when a TCDS is preferred over a CDS [2]. However, THP is not reliable when the network topology changes frequently, and there is a clear trade-off between reliability and efficiency.

To improve the route discovery process of on demand routing protocols, and to address the lack of reliability in the presence of high mobility, we presented two enhancements to THP. The new algorithm was called *Three-hop Horizon Enhanced Pruning* (THEP).

A *virtual radio range* (VR), shorter than the physical radio range (RR), is used for gathering information about the two-hop neighborhood. Instead of using two different transmission powers, which can increase interference, we use a single radio range while still managing to have a *buffer zone* in which nodes can move without compromising network connectivity.

Upon receiving a broadcast packet, the forwarder list in the packet header is analyzed together with the current information about the local neighborhood. This is done to find inconsistencies between the most up-to-date *one-hop dominating list* and the one used by the sender to compute the forwarder list. Changes in the local topology may have impacted the *one-hop dominating list*. If that is the case, a node may decide to relay a broadcast packet even though the node was not selected as a forwarder by the sender.

We have implemented the two enhancements to THP in AODV (the resulting AODV modifications is called AODV-THEP) using the *Qualnet* simulator. Extensive simulation results show that AODV-THEP attains better performance than AODV in terms of delivery ratio, control overhead, packet collision, and end-to-end delay.

REFERENCES

- [1] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," *Wireless Networks*, vol. 8, pp. 153–167, 2002.
- [2] M. A. Spohn and J. J. Garcia-Luna-Aceves, "Improving broadcast operations in ad hoc networks using two-hop connected dominating sets," in *Proceedings of The IEEE GLOBECOM Wireless Ad hoc and Sensor Networks*, november 2004.
- [3] C. Perkins, "Ad-hoc on-demand distance vector routing," in *Second IEEE Workshop on Mobile Computing Systems and Applications*, 1999.
- [4] J. Wu and F. Dai, "Mobility management and its applications in efficient broadcasting in mobile ad hoc networks," in *IEEE INFOCOM*, 2004.
- [5] B. Williams and T. Camp, "Comparison of broadcasting techniques for mobile ad hoc networks," in *Proceedings of MOBIHOC*, 2002, pp. 194–205.
- [6] M. R. Garey and D. S. Johnson, *Computers and Intractability*. Freeman, San Francisco., 1978.
- [7] P.-J. Wan, K. M. Alzoubi, and O. Frieder, "Distributed construction of connected dominating set in wireless ad hoc networks," in *Proceedings of INFOCOM 2002*, June 2002.
- [8] H. Lim and C. Kim, "Flooding in wireless ad hoc networks," *Computer Communications*, vol. 24, february 2001.
- [9] W. Lou and J. Wu, "On reducing broadcast redundancy in ad hoc wireless networks," *IEEE Transactions on Mobile Computing*, vol. 1, no. 2, April-June 2002.
- [10] M. A. Spohn and J. J. Garcia-Luna-Aceves, "Enhanced dominant pruning applied to the route discovery process of on-demand routing protocols," in *Proceedings of the 12th IEEE ICCCN*, october 2003.
- [11] J. Wu and F. Dai, "Mobility-sensitive topology control in mobile ad hoc networks," in *IEEE IPDPS*, 2004.
- [12] J. Hightower and G. Borriello, "A survey and taxonomy of location systems for ubiquitous computing," University of Washington, Computer Science and Engineering, Tech. Rep. UW-CSE 01-08-03, 2001.
- [13] "Scalable network technologies," 2003, qualnet simulator. <http://www.scalable-networks.com>.
- [14] Y. Wang and J. J. Garcia-Luna-Aceves, "Spatial Reuse and Collision Avoidance in Ad Hoc Networks with Directional Antennas," in *Proceedings of GLOBECOM 2002*, Taipei, Taiwan, November 2002.